

REMARKSI. Introduction

In response to the Office Action dated October 3, 2002, claims 1, 9, and 17 have been amended. Claims 1-4, 6-12, 14-20, and 22-24 remain in the application. Re-examination and re-consideration of the application, as amended, is requested.

II. Claim Amendments

Applicants' attorney has made amendments to the claims as indicated above. These amendments were made solely for the purpose of clarifying the language of the claims, and were not required for patentability or to distinguish the claims over the prior art.

III. Non-Art Rejections

In paragraphs (3)-(4) of the Office Action, claims 1-4, 6-12, 14-20, and 22-24 were rejected under 35 U.S.C. §112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which Applicants regard as the invention. Specifically, the Office Action has rejected claims 1-4, 6-12, 14-20, and 22-24 as follows:

Claims 1, 9 and 17 cite the feature of combining unlocking and writing. It is not clear from the specification or the claim what combining the steps mean. Are they performed sequentially or concurrently.

Claims 4, 12 and 20 cite the feature of combining reading and locking. Again it is not clear from the specification or the claim what combining the steps mean. Are they performed sequentially or concurrently.

The Examiner has interpreted these features to mean that the steps are performed sequentially.

The specification provides that in an enhanced protocol, the unlocking and writing steps are combined (see specification page 10, lines 6-14, and page 11, line 23 to page 12, line 1). With this enhanced protocol, the overhead for ensuring accuracy is minimized, flexible, and does not increase with additional nodes (see page 10, lines 6-12, and page 12, lines 1-2). In this regard, the advantages lie in reducing the overhead (i.e., the number of commands transmitted). The specification describes the use of the RESERVE and RELEASE commands to accomplish the locking and unlocking steps (see page 10, line 16 to page 11, line 16). The original dependent claims also provided for issuing commands for the locking and unlocking. Accordingly, commands are utilized to perform locking and unlocking. The amended claims and specification provide for combining the

writing and unlocking steps (i.e., commands). In this regard, the claims have been amended to clarify that the writing and unlocking steps are combined into a single command.

Accordingly, Applicants have amended the claims to provide that the writing and unlocking steps are combined into a single command. Thus, regardless of whether the steps are performed sequentially or concurrently, the commands for performing the steps are combined so that only a single command is issued. The use of a single command that performs both unlocking and writing reduces the overhead for ensuring accuracy.

In view of the above, Applicants submit that the rejection under 35 U.S.C. §112 is now moot and the claims are in allowable form.

IV. Prior Art Rejections of Claims 6, 8, 14, 16, 22, and 24

In paragraph (8) of the Office Action, claims 6, 14, and 22 were rejected under 35 U.S.C. 103(a) as being unpatentable over Hodges, U.S. Patent No. 6,219,751 B1 (Hodges) in view of IBM Technical Disclosure Bulletin, "Limited Distributed DASD Checksum" (IBM). In paragraph (9) of the Office Action, claims 8, 16, and 24 were rejected under 35 U.S.C. §103(a) as being unpatentable over Hodges in view of Lyons, U.S. Patent No. 6,101,615 (Lyons).

In response to the prior Office Action, Applicants submitted a Statement of Common Ownership with respect to the Hodges references. Accordingly, Hodges cannot be utilized in a rejection under 35 U.S.C. 103(a). The rejection of claims 6, 8, 14, 16, 22, and 24 continue to utilize Hodges despite the submission of the Statement of Common Ownership. Accordingly, Applicants submit that the rejection of claims 6, 8, 14, 16, 22, and 24 should be withdrawn in view of the previously submitted Statement of Common Ownership and that these claims are now in allowable form.

V. Prior Art Rejections

In paragraphs (5)-(6) of the Office Action, claims 1, 4, 7, 9, 12, 15, 17, 20, and 23 were rejected under 35 U.S.C. §102(b) as being anticipated by Menon et al., U.S. Patent No. 5,574,882 (Menon). In paragraph (10) of the Office Action, claims 2-3, 10-11, and 18-19 were rejected under 35 U.S.C. §103(a) as being unpatentable over Menon in view of Ofer, U.S. Patent No. 5,892,955 (Ofer).

Independent claims 1, 9, and 17 are generally directed to updating parity data in a RAID clustered environment. Specifically, parity is locked, read, and new parity data is generated. The new parity data is written and then unlocked. Further, the commands for writing and unlocking the parity are combined into a single command. As described above, the combining of these commands into a single command provides for less overhead than other methods for any number of nodes in a cluster.

The cited references do not teach nor suggest these various elements of Applicants' independent claims.

Menon merely describes a system and method that is used by software implemented Redundancy Array of Inexpensive Disk (RAID) arrays to achieve adequate performance and reliability, as well as to improve performance or low cost hardware Raids. The enhancements to the basic RAID implementation speeds up recovery time for software RAIDS. A method is provided for storing data in an array of storage devices. A plurality of block locations on the storage devices are logically arranged as a parity group wherein a parity block stored in a block location as part of a parity group is logically derived from the combination of data blocks stored in the parity group, and each block in a parity group is stored on a different storage device. A plurality of parity groups are grouped into a parity group set. A request is received to write a new data block location on a storage device. The old data block stored at the block location is read. The new data block is written to the block location. When the parity set is in an unmodified state prior to the current write, an indicator is written to the storage device that the parity group set is in a modified state. In a preferred embodiment, this enhancement uses a bit map stored on disk, called Parity Group Set, (PGS) bit map, to mark inconsistent parity groups, replacing the Non-Volatile Random Access Memory, (NVRAM) used for similar purposes by hardware RAIDs. Further enhancements optimized sequential input/output, (I/O) data stream.

The cited portions of Menon (col. 6, lines 45-46) lack any description whatsoever of combining the writing and unlocking commands into a single command. Instead, Menon merely describe a parallel write method that allows the writing of data and parity to proceed in parallel so that locks are held for a shorter period of time (see col. 6, lines 2-5). Menon's parallelism is achieved by affixing time stamps to disk blocks (see col. 6, lines 5-6) and not by combining commands. Once the data is safely written to disk, the parity group is unlocked (see col. 6, lines 43-46). Thus, Menon

merely describes parallel writing to disk and parity. Thereafter, once complete, the parity is unlocked. There is no discussion, explicit or implicit, in Menon, that describes combining the writing and unlocking (and not locking) of parity. Further, the claims provide for combining the commands into a single command. Menon fails to describe any such command combining as claimed.

Moreover, the various elements of Applicants' claimed invention together provide operational advantages over Menon, Hodges, IBM, Lyons, and Ofer. In addition, Applicants' invention solves problems not recognized by Menon, Hodges, IBM, Lyons, and Ofer.

Thus, Applicants submit that independent claims 1, 9, and 17 are allowable over Menon, Hodges, IBM, Lyons, and Ofer. Further, dependent claims 2-8, 10-16, and 18-24 are submitted to be allowable over Menon, Hodges, IBM, Lyons, and Ofer in the same manner, because they are dependent on independent claims 1, 9, and 17, respectively, and thus contain all the limitations of the independent claims. In addition, dependent claims 2-8, 10-16, and 18-24 recite additional novel elements not shown by Menon, Hodges, IBM, Lyons, and Ofer.

VI. Conclusion

In view of the above, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited. Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicants' undersigned attorney.

Respectfully submitted,

GATES & COOPER LLP
Attorneys for Applicant(s)

Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, California 90045
(310) 641-8797

Date: December 30, 2002

By: G.H.Gates
Name: George H. Gates
Reg. No.: 33,500

GHG/JSF/io

G&C 30879.80-US-01

APPENDIX: VERSION WITH MARKINGS TO SHOW CHANGES MADE

1. (TWICE AMENDED) A method of updating parity data in a redundant array of independent disk (RAID) clustered environment comprising:

- (a) locking parity data, without communicating with other nodes, for data managed in SCSI (small computer systems interface) disks in a RAID clustered system, wherein the locking prevents other nodes from modifying the parity data;
- (b) reading the parity data;
- (c) generating new parity data by exclusive oring data from a first node and a second node;
- (d) writing the parity data to a SCSI disk in the RAID clustered system; and
- (e) unlocking the parity data, wherein [the] commands for writing and unlocking [steps] are combined into a single command.

9. (TWICE AMENDED) An apparatus for updating parity data in a redundant array of independent disk (RAID) clustered environment comprising:

- (a) a plurality of SCSI (small computer systems interface) storage devices organized in a RAID clustered system;
- (b) data stored in the plurality of SCSI storage devices;
- (c) a first node, operatively coupled to the plurality of SCSI storage devices, that manages storage and retrieval of the data in the plurality of SCSI storage devices, wherein the first node is configured to:

- (i) lock parity data without communicating with other nodes, wherein a lock prevents other nodes from modifying the parity data;
- (ii) read the parity data;
- (iii) generate new parity data by exclusive oring data from two nodes;
- (iv) write the parity data to a SCSI disk in the RAID clustered system; and
- (v) unlock the parity data, wherein logic for writing and unlocking are combined into a single command.

17. (TWICE AMENDED) An article of manufacture, embodying logic to perform method steps of updating parity data in a redundant array of independent disk (RAID) clustered environment, the method steps comprising the steps of:

- (a) locking parity data without communicating with other nodes, wherein the locking prevents other nodes from modifying the parity data;
- (b) reading the parity data;
- (c) generating new parity data by exclusive oring data from two nodes;
- (d) writing the parity data to a SCSI (small computer systems interface) disk in the RAID clustered system; and
- (e) unlocking the parity data, wherein [the]commands for writing and unlocking [steps] are combined into a single command.